


amdaahl

580 Technical Introduction



The Amdahl 580: An Evolutionary Extension of the Amdahl Concept

The new Amdahl 580 computer system provides an average of two times the processing power of the Amdahl 470V/8 in typical commercial environments. This evolutionary extension to the Amdahl product line is made possible by improvements in:


- Large Scale Integration (LSI) circuit density, packaging, and power requirements.
- High-speed Random Access Memories (RAMs) for distributed control storage and High-Speed Buffers (HSBs).
- System architecture and design, including microcode and a new concept—Macrocode.

The Amdahl 580 concept is consistent with Amdahl's fundamental objective of applying advanced technology and innovative design to large-scale computers providing the following principal benefits:

- **COMPATIBILITY** with IBM large-scale and Amdahl 470 series computer systems, and enhanced **FLEXIBILITY** to maintain future compatibility to preserve the billions of dollars of investment in application software, employee training, and installed hardware.
- Improved **PERFORMANCE** and **PRICE/PERFORMANCE** to reduce the hardware investment required to meet accelerating application development and growth while improving productivity and reducing operating costs.
- Improved **AVAILABILITY** to increase the return on the computer investment by providing more usable hours of computer processing.

The Amdahl 580's many innovations and features will be presented according to their impact on the principal benefits of **COMPATIBILITY, PERFORMANCE, and AVAILABILITY**.


Amdahl's 470 computer systems have been enthusiastically accepted by a worldwide community of large-scale, general purpose computer users. The Amdahl 580 continues in the 470 tradition: it preserves the investment of Amdahl customers and prospects by providing an upward compatible growth path.



470 Multiple Chip Carrier (MCC), holds up to 42 chips.


A new Multiple Chip Carrier (MCC) accommodates 121 chips.

A new, 400 picosecond ECL/LSI chip packs 400 circuits per chip.




Unmounted 470 ECL/LSI chip has 25 cells with up to four circuits each.

470 LSI chip carrier has three fins for air cooling.



A modified chip carrier permits the continued use of Amdahl air cooling.





60 MCCs are planar mounted in 470V/8.

A new stack design contains almost all of the logic for the basic Amdahl 580 on eight MCCs and occupies only 5.6 cubic feet of space.

Amdahl 470 mainframe.


A new basic mainframe design measures only ten feet in length and occupies only 33 square feet of floor space.






A new Multiple Chip Carrier (MCC) accommodates 121 chips.

A new, 400 picosecond ECL/LSI chip packs 400 circuits per chip.




A modified chip carrier permits the continued use of Amdahl air cooling.



A new stack design contains almost all of the logic for the basic Amdahl 580 on eight MCCs and occupies only 5.6 cubic feet of space.

A new basic mainframe design measures only ten feet in length and occupies only 33 square feet of floor space.



A New Standard in Computing Price/Performance

Amdahl 580 Features Summary

The Amdahl 580 is a fully compatible extension of the Amdahl 470 series. Through new advancements in technology and system design, it delivers an average of two times the power of the Amdahl 470V/8 to the typical commercial user. This improvement in performance is largely due to new design concepts that increase the maximum instruction rate to one per machine cycle and significantly decrease the cycles required to complete the average instruction.

Up to 32 megabytes of Main Memory and dual 32K High-Speed Buffers (HSBs) utilizing a two-way, set-associative design are also included for efficient and flexible memory access.

In addition, there can be up to 34 Input/Output (I/O) channels, including any combination of up to 32 Block Multiplexer Channels and up to two Byte Multiplexer Channels. The I/O channels, together with a dual bus communications structure, create a fast and efficient I/O system capable of handling data rates up to six megabytes per second per channel, and up to 50 megabytes per second in the aggregate.

The Amdahl 580 system is designed to expand beyond 32 megabytes and to allow additional processing power to be attached to the basic system in the future.

The Amdahl 580's Console provides full operator control from up to four CRT/keyboard stations. The Console contains a powerful processor of its own, which executes a subset of the Amdahl 580 system instructions and provides complete local or remote diagnostic and maintenance functions.

Outstanding PERFORMANCE through Advanced System Design

The Amdahl 580's design improvements include a five-phase, instruction-per-cycle pipeline that permits overlapping the execution of five instructions at a time. As an instruction moves through the pipeline, another can immediately follow. The result is a maximum rate of one instruction processed per machine cycle. Pipeline utilization is more efficient due to significant improvements in functional unit organization and instruction algorithm design.

The Amdahl 580's physical organization permits the grouping of an entire function, such as instruction execution, on a single Multiple Chip Carrier (MCC). Significant performance improvements result from the shortened data paths made possible by less inter-chip and inter-MCC data traffic.

The dual bus structure of the Amdahl 580 efficiently transfers data between functional units on eight-byte wide data paths. I/O traffic and CPU-to-Main Memory traffic move with a minimum of interference. The bus organization, together with the use of dual independent High-Speed Buffers (HSBs) for instruction and operand data, improve performance by reducing the instances in which the CPU must wait for needed data.

Continuing COMPATIBILITY through Flexible Architecture

The Amdahl 580 system is compatible with IBM's large-scale computer systems and the Amdahl 470 series. Software developed for any of these systems will run on the Amdahl 580 with no modification, except in those limited cases where the software is model dependent. More important, the Amdahl 580 design and technology are adaptable to future changes to maintain compatibility.

The Amdahl 580 design allows four different techniques with which to retain compatibility:

(1) hardware modification, (2) microcode, (3) Macrocode, or (4) software emulation.

The Amdahl 580 makes extensive use of Distributed Microcode. It incorporates a new high-speed Random Access Memory (RAM) chip to distribute microcode physically to the functional units where it is used, without incurring a performance degradation. The microcode structure for each unit is tailored to fit its particular function.

Macrocode, a new class of firmware, is an Amdahl innovation. Macrocode provides additional flexibility to accommodate a wide range of new functions. Macrocode is more flexible than microcode, since it does not share the same complex relationship to the hardware. New functions can be implemented in Macrocode more easily than in microcode.

The Amdahl 580 provides continued peripheral device and channel compatibility. The I/O Controller (IOC) in the I/O Processors (IOPs) is microcoded for flexibility. The Amdahl 580's Channel Interface Handler logic is modular and independent of other system logic. If a new I/O protocol is required by the market, redesign is limited to a single printed circuit card.

Improved AVAILABILITY through Fewer Components

The more powerful a computer, the more critical is its availability—the time spent performing useful work. The Amdahl 580 provides improved availability through increased reliability and serviceability.

Reliability has been improved through a reduction in the number of components and interconnections—a reduction made possible by denser chips and higher capacity MCCs. Inter-MCC signals are carried by printed circuit boards, rather than by discrete wiring. Error Checking and Correction (ECC) or parity checks are used on all data paths, and residue and parity checks are used in arithmetic units.

The Amdahl 580 sets high standards in serviceability and Mean-Time-To-Repair (MTTR). Increased circuit density permits an entire functional unit to reside on a single MCC, resulting in fast fault isolation and speedy repair.

A fault isolation strategy enables easier identification of a failing MCC. Each MCC contains logic which records the identity of the circuit on that MCC detecting an error. In many cases it also records the sources of data and control signals and latch contents. This information is used to decide which Field Replaceable Unit (FRU) to exchange. Just knowing which functional unit is associated with the problem symptoms allows selection, in many cases, of the correct FRU for replacement. Diagnostics can be performed over telephone lines anywhere in the world by Amdahl specialists at one of several Amdahl Diagnostic Assistance Center (AMDAC) sites.

The Amdahl 580 includes these system enhancements to facilitate rapid diagnosis:

- Built-in logic scan facilities to determine and/or set the contents of system latches.
- Console and AMDAC facilities.
- Microcode and bus-transaction history RAMs for fault tracing.
- History RAM to record Main Memory correctable errors.
- A separate console support processor to assist in diagnosing Console failures.

Amdahl 580 Technological Innovations


The Amdahl 580's improved reliability and compact size derive from the use of dense Large Scale Integration (LSI) chips utilizing Emitter-Coupled Logic (ECL) with gate delays on the order of 400 picoseconds (trillionths of a second). The Amdahl 580 chip can contain up to 400 of these high-speed circuits. A new, more efficient process technology results in a faster circuit which generates only one quarter of the heat as the same circuit on an Amdahl 470 chip. Because of greater circuit utilization, the Amdahl 580 chip generates only slightly more

heat, but a modified chip carrier easily dissipates this additional heat with air cooling.

A new high-speed LSI Random Access Memory (RAM) chip has been developed for buffer storage, registers, and microcode on the same MCC with logic chips. This new RAM makes Distributed Microcode and one-cycle High-Speed Buffers (HSBs) possible.



The Multiple Chip Carrier (MCC) in the Amdahl 580 can hold up to 121 chips and implements an entire system function. The number of layers in the Amdahl 580 MCC has been increased to 14 to accommodate more internal data paths. Eight MCCs implement almost all the logic for the basic Amdahl 580 and are arranged in a

stack which is about one-fourteenth the size of the Amdahl 470 LSI circuitry (5.6 cubic feet vs. 79 cubic feet). A ninth MCC fits in the same stack space when a second Input/Output Processor (IOP) is required to expand the system beyond the basic channel configuration. The stack design incorporates 12-layer printed circuit boards for MCC-to-MCC interconnections, thus reducing signal path lengths and improving system reliability.




MCC Stack

400-Circuit Chip



New RAM



121-Chip MCC



System Overview

The Amdahl 580 system consists of a mainframe, a Power Distribution Unit (PDU), and up to four operators' Console CRT/keyboard units.

Mainframe

The Amdahl 580 mainframe contains the stack, the Main Storage Unit (MSU) which provides up to 32 megabytes of Main Memory, and a frame housing the Channel Interface Handler cards and certain Console components.

Stack

The stack for the basic Amdahl 580 holds eight functional units—six implemented on their own Multiple Chip Carriers (MCCs) and two buffer units distributed across two MCCs. The second Input/Output Processor (IOP), a ninth functional unit and MCC, may be optionally included in the stack. The MCCs are mounted horizontally in the stack, and most are connected by two uni-

directional communication buses—the A-Bus and the B-Bus. The buses are distributed by the printed circuit boards that form the sides of the stack.

Five of the stack-implemented functional units comprise the Amdahl 580 Central Processing Unit (CPU):

- Instruction Unit (I-Unit): Fetches, decodes, and controls instructions and controls the CPU.
- Execution Unit (E-Unit): Provides computational facilities.
- Storage Unit (S-Unit): Controls the Amdahl 580's instruction operand storage and retrieval facilities.
- Instruction Buffer (I-Buffer): Provides High-Speed Buffer storage for instruction streams.
- Operand Buffer (O-Buffer): Provides High-Speed Buffer storage for operand data.

Operation of the functional units within the CPU is overlapped,

or pipelined. This pipeline organization allows up to five instructions to be in some phase of execution simultaneously.

The three remaining functional units contained in the basic Amdahl 580 stack are:

- Input/Output Processor (IOP): Receives and processes I/O requests from the CPU and provides up to 16 Block Multiplexer Channels.*
- Time-sliced Console Processor: Communicates with the CPU to provide system control and up to two Byte Multiplexer Channels.
- Memory Bus Controller (MBC): Provides Main Memory and bus control, system-wide coordination functions, and timing facilities.


*A second IOP may be included to provide up to 16 additional Block Multiplexer Channels.

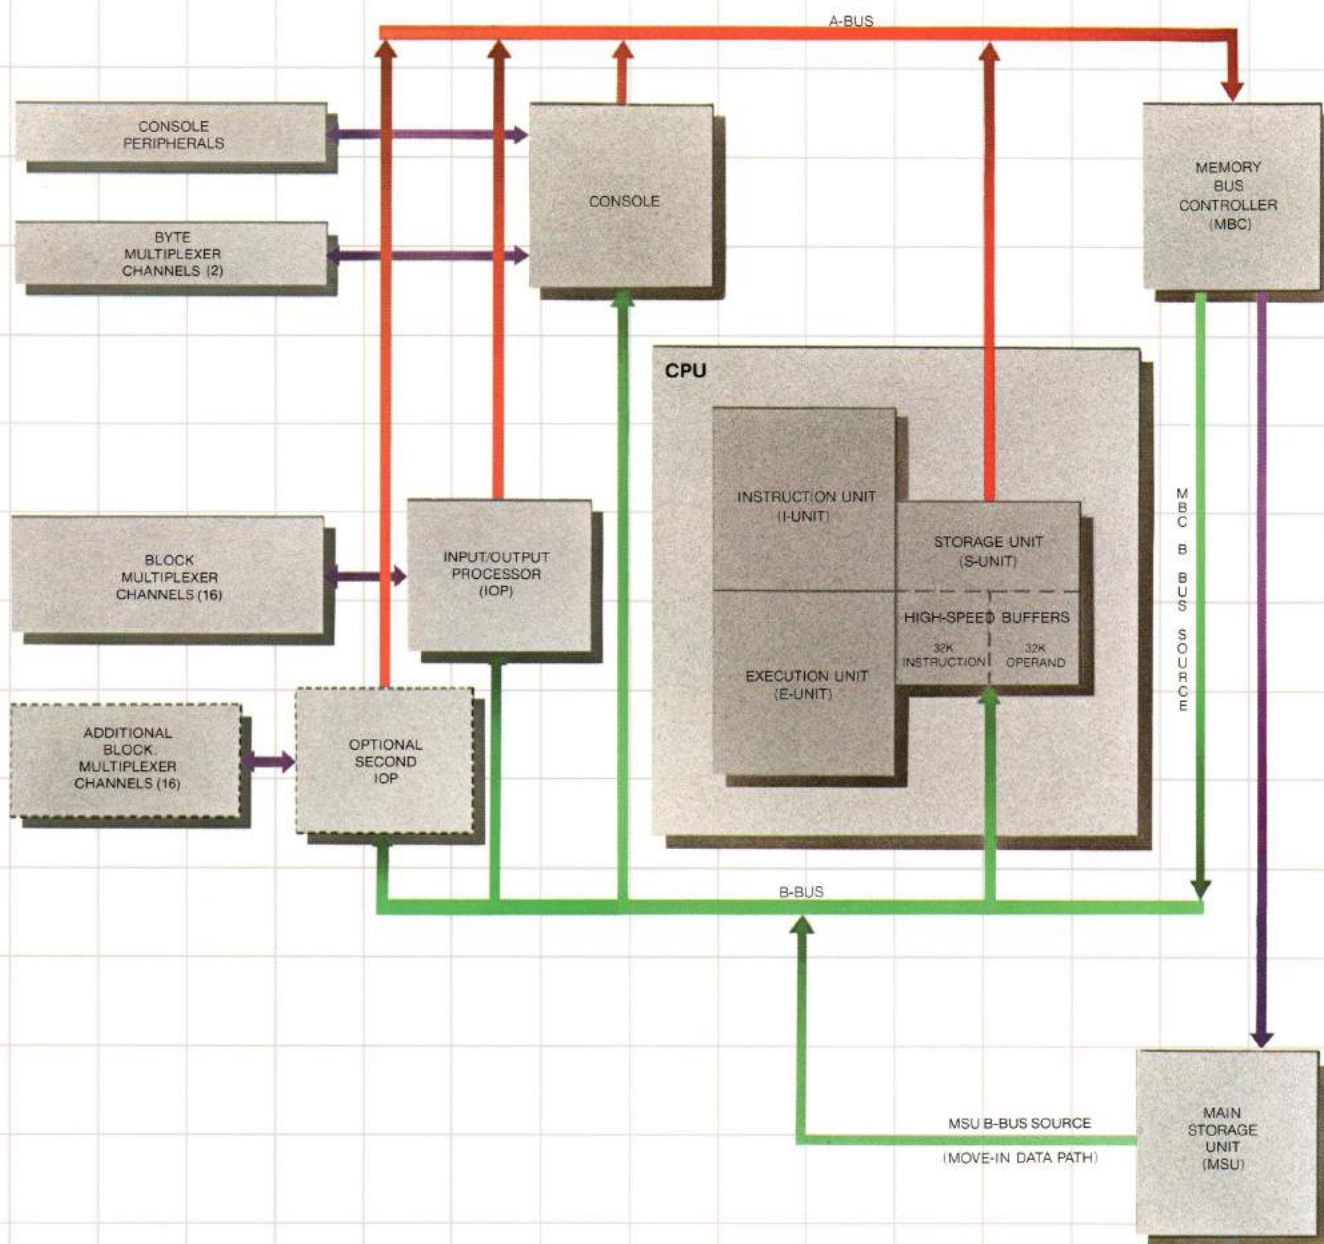
Bus Structure Simplifies Data Flow

The Amdahl 580's dual bus structure provides data paths which have a common interface with each functional unit. Each bus is unidirectional with a data path 72 bits wide to carry a 64-bit (eight-byte) message plus byte parity information. The A-Bus carries data from the Console, IOP, and CPU to the MBC. The B-Bus carries data from the MBC or MSU to the Console, IOP, and CPU.

The Amdahl 580's dual bus structure is a disciplined approach to connecting the functional units. Because the bus data paths are integral parts of the stack walls, the resulting path lengths are shorter, physical connections are simplified, and connections among functional units are minimized. This design contributes significantly to the Amdahl 580's performance and reliability.

The Amdahl 580 Mainframe





AMDAHL 580 SYSTEM OVERVIEW

CPU Organization

In the Amdahl 580 CPU, two sets of functions are continuously performed in parallel:

- Instruction Fetch (I-Fetch) and
- Instruction Execution.

The diagram at the right illustrates how these processes are carried out.

Instruction Fetch

The I-Fetch process provides a double word of instruction stream every cycle and holds the instructions in the Instruction Unit (I-Unit) for use whenever the execution process is ready. The I-Unit's I-Fetch mechanism controls this process and uses the I-Unit's Instruction Address Generator (IAG), the Instruction Buffer (I-Buffer), and the Storage Unit (S-Unit) to fetch instructions which are then held in its own Instruction Word Buffer (IWB), the interface between I-Fetch and pipeline control.

The IWB can hold four half words of instruction data. The top two half words contain instruction data entering the first phase of the execution process, and the bottom two half words contain instructions awaiting execution. Typically, at the end of a cycle, the instruction completing the first execution phase is bubbled-up or pushed out of the top one or two half words of the IWB; the instructions in the bottom two or three half words are bubbled-up, and previously requested instruction data arriving from the I-Buffer is inserted in the now vacant bottom half words.

In each cycle, the I-Fetch mechanism uses the IAG to calculate the address of the instruction data that will be needed in the next cycle to fill the IWB. That address is sent to the I-Buffer and S-Unit, which access the instruction data and return it to the I-Fetch mechanism in the next cycle.

In the case of a branch instruction, the operand fetch mechanism requests an I-Buffer access of the target instruction stream. If the branch is taken, instruction data from the target fetch fills the IWB, causing only a single cycle delay before the target instruction begins execution. Otherwise, the target fetch is cancelled, and the next sequential instruction is bubbled-up and is ready for immediate execution. In many cases, the Amdahl 580's short five-phase pipeline allows an early decision on which instruction stream to use and makes extensive buffering of target instruction streams unnecessary.

Instruction Execution

At the maximum execution rate, instructions held in the IWB are presented to the first phase of the execution process, or pipeline, every cycle. There are five phases of pipeline flow for each instruction, and at the maximum rate each instruction advances one phase per cycle:

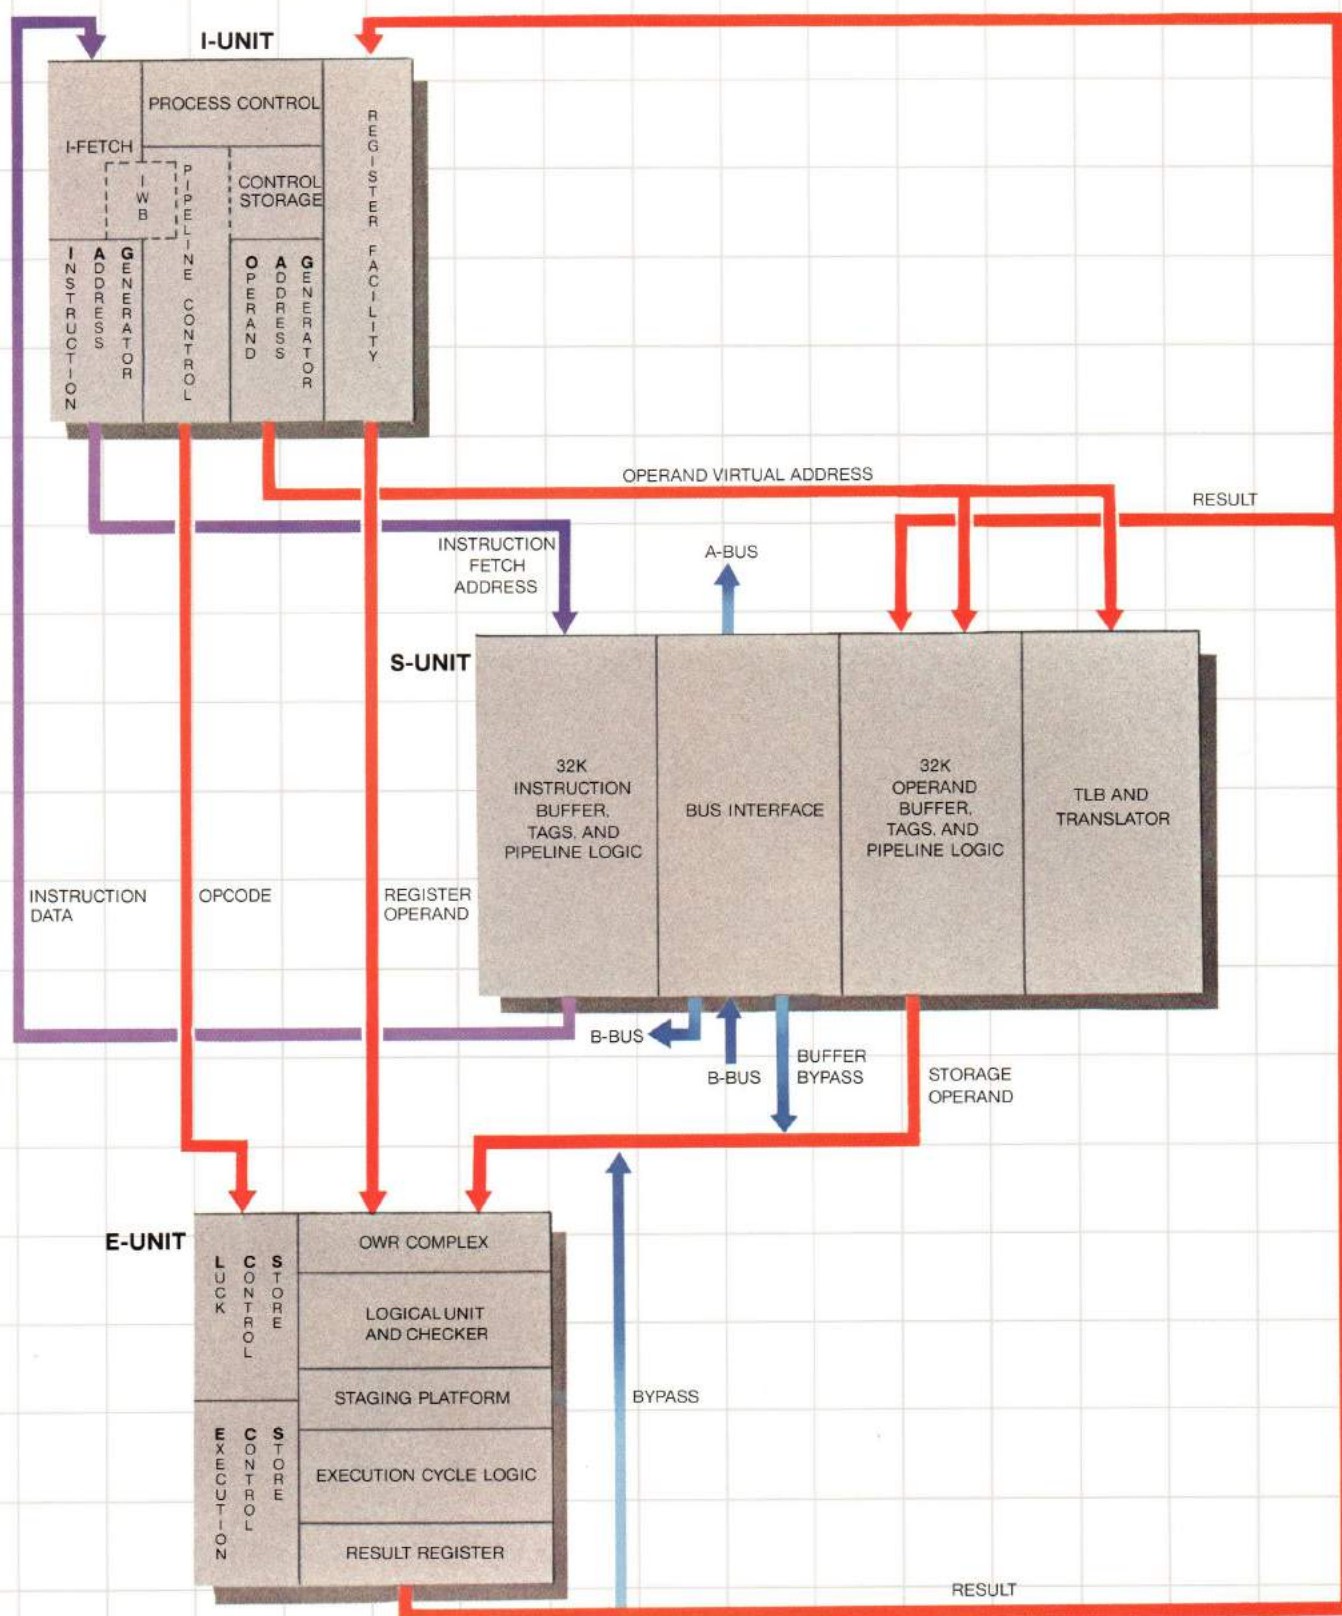
- Generate (G) Phase: The instruction at the top of the IWB is decoded; the opcode is checked for validity; the operand virtual address is calculated by the Operand Address Generator (OAG) and sent to the Operand Buffer (O-Buffer) and S-Unit; the pipeline controls are set for later phases of instruction execution; the opcode is sent to the Execution Unit (E-Unit); the I-Unit Control Store (ICS) is accessed for further I-Unit processing if necessary.

- Buffer (B) Phase: The O-Buffer is accessed; the I-Unit accesses register operands; the control word for the next execution phase is accessed using the opcode as an index into the Logical Unit and Checker (LUCK) Control Store (LCS); the operand data from the registers or the O-Buffer is gated into the E-Unit Operand Word Register (OWR) complex.
- LUCK (L) Phase: The LUCK processes the operands as specified by the control word from LCS and sets condition codes if appropriate; logical functions, comparison, and certain data manipulations are done here; the Execution Control Store (ECS) is accessed to obtain the control word needed for the next cycle.
- Execution (E) Phase: The calculations specified by the ECS control word are executed, and the result is placed in the result register.
- Write (W) Phase: The result is stored in the I-Unit register facility or O-Buffer as appropriate.



CPU Organization

- Separate Instruction Fetch and Instruction Execution are continuously performed in parallel.



CENTRAL PROCESSING UNIT (CPU)

Instruction Unit (I-Unit)

The I-Unit controls the execution of instructions and the processing of interruptions within the Amdahl 580. The major functions performed by the I-Unit are:

- Fetching, buffering, and decoding instructions.
- Calculating effective addresses for operand and instruction fetches.
- Providing access to the register file for operands and addressing purposes.
- Controlling the machine state and processing all interrupts and machine checks.
- Controlling the Storage Unit (S-Unit), Execution Unit (E-Unit), and Input/Output Processors (IOPs) to accomplish overlapped pipeline execution.

Performance

The Amdahl 580's instruction-per-cycle pipeline enhances performance by improving instruction throughput, especially on most frequently used instructions, such as loads and stores. With the short five-phase pipeline, condition codes from the most common branch predecessor instructions—such as Load and Test Register, Compare Logical Immediate, and Test Under Mask—are available in the third phase and cause no delays in the branch decision. Condition codes not set by the execution results of other branch predecessor instructions are available in the fourth phase and cause only a one-cycle hole in the pipeline. Pipeline interlocks caused by instruction stream dependencies are reduced and optimized for frequently encountered instruction sequences.

The instruction fetch mechanism is independent of the execution pipeline logic to reduce interference. A double word of instruction information can be fetched every cycle to keep the instruction buffers

full. The Amdahl 580's instruction buffering technique minimizes pipeline holes in the case of branching, yet avoids fetching excessive numbers of never-to-be executed instructions.

In order to have the greatest effect on overall performance, improvements in I-Unit algorithms are focused on instructions which use the most cycles in typical customer environments. The net result of the Amdahl 580's design is an average instruction rate twice that of the Amdahl 470V/8 system in typical commercial environments.

Compatibility

The Amdahl 580 I-Unit is compatible with System/370 Principles of Operation opcodes and implements these instructions with an optimized mixture of hardware, microcode, and Macrocode.

Crucial functions, either common to all opcodes or requiring extensive decision resolution, are implemented in hardware to ensure the fastest possible execution. Examples include register conflict resolution and operand effective address generation. Other functions can utilize the highly efficient microcode execution facility located on the same MCC as the I-Unit and tailored to the I-Unit's needs. Microcode provides the capability to modify or add a limited number of new functions.

Availability

I-Unit reliability is enhanced by instruction retry. I-Unit serviceability elements include two last-branch address registers for easier instruction tracing and a microcode history RAM for fault tracing.

The diagram illustrates the Amdahl 580 five-phase pipeline overlap. Once the I-Fetch mechanism fetches the instruction, pipeline logic takes over, and instruction number 1 enters the Generate phase. At the end of the machine cycle this phase is completed; instruction number 1 progresses to the Buffer phase; and instruction number 2 enters the pipeline in the Generate phase.

Typically, by the fifth machine cycle, five instructions are in the pipeline, each executing in a different phase.



I-Unit

- Five-phase pipeline organization allows one instruction per cycle at the maximum execution rate, and five instructions to be in execution phases simultaneously.
- Early condition code setting and reduced pipeline interlocks improve instruction throughput.
- Macrocode, a new class of firmware, offers flexibility for future extension.
- Microcode history RAM and instruction retry improve availability.

AMDAHL 580 PIPELINE FLOW

INSTRUCTION FLOW

PIPELINE PHASE	FUNCTION
GENERATE	CHECK OPCODE VALIDITY GENERATE OPERAND ADDRESS SEND OPERAND FETCH REQUEST TO S-UNIT ACCESS I-UNIT CONTROL STORE (ICS) SEND OPCODE TO E-UNIT
BUFFER	ACCESS OPERAND BUFFER ACCESS OPERAND IN REGISTER ACCESS LUCK CONTROL STORE (LCS) INGATE OPERAND DATA TO OWR
LUCK	PERFORM FUNCTIONS SPECIFIED BY LCS ACCESS EXECUTION CONTROL STORE (ECS)
EXECUTION	PERFORM FUNCTIONS SPECIFIED BY ECS
WRITE	WRITE RESULTS TO REGISTER OR BUFFER

AMDAHL 580 PIPELINE OVERLAP

G B L E W	1	2	3	4	5	6	7	8	9	10
		1	2	3	4	5	6	7	8	9
			1	2	3	4	5	6	7	8
				1	2	3	4	5	6	7
					1	2	3	4	5	6
PIPELINE PHASE	1	2	3	4	5	6	7	8	9	10
MACHINE CYCLES ▶▶▶										

Storage Unit (S-Unit) and High-Speed Buffers (HSBs)

The S-Unit, together with the 32K Instruction Buffer (I-Buffer) and 32K Operand Buffer (O-Buffer), provide High-Speed Buffer storage for instructions and operands. With only occasional access to the relatively slower Main Memory, they effectively keep the CPU operating close to its maximum instruction rate.

The S-Unit receives and processes all Instruction Unit (I-Unit) data requests. It translates virtual addresses to absolute addresses and maintains the Translation Lookaside Buffer (TLB) to provide quick virtual-to-absolute translations. The S-Unit also controls data traffic between the CPU data buffers and Main Memory and provides the bus interface between the CPU and the rest of the Amdahl 580.

Performance

The S-Unit and buffers are organized to access a double word of data from each buffer on every cycle by simultaneously accessing the four storage arrays which contain data lines, tags, and the TLB.

Separate I-Fetch and execution processes make it necessary for the S-Unit to have separate instruction and operand fetch pipelines, buffers, and tags. Both buffers are two-way, set-associative and are organized into primary and alternate halves of 512 32-byte lines to reduce the incidence of missing buffer lines.

If the S-Unit instruction and operand pipelines both had to access the TLB for virtual-to-absolute translation information, the resulting contention would degrade performance. To avoid this problem, the S-Unit instruction pipeline uses a special register to maintain translation information for the current executing page instead of accessing the TLB.

In the event that a line of requested data is not in the buffer, the S-Unit sends a message over the A-Bus to the Main Storage Controller (MSC) in the Memory Bus Controller (MBC), requesting that the line be accessed in Main Memory and returned over the Move-in Data Path. A line of data is moved out of the buffer to Main Memory by sending a message with a header segment and four quarterline data segments to the MSC over the A-Bus and then on to Main Memory.

The 512-entry TLB is also organized into primary and alternate halves of 256 translations each to speed access to virtual-to-absolute address translations. Segment Table Origin (STO) information is included in each TLB entry to eliminate the need for a STO stack and to provide faster access.

Availability

The Amdahl 580's Main Memory contains Error Checking and Correction (ECC) information to permit correction of single-bit errors and detection of all double-bit errors during move-ins. The Operand Buffer contains ECC information to permit the same correction and detection of errors during move-outs. Instruction Buffer data may be refetched by the Recovery Management System (RMS) from Main Memory to correct errors when they are detected.

An Operand High-Speed Buffer access request from the I-Unit consists of an operand virtual address and Length, Justification, and Rotation (LJR) information. The access request is honored by simultaneously accessing four storage arrays, performing two comparisons, and sending the output to the E-Unit Operand Word Register (OWR) complex.

The four storage arrays are the Data Array, the Tag Array, the Data Select Array, and the Translation Lookaside Buffer (TLB) Array. The Data Array contains the actual data lines and is organized into primary and alternate halves of 512 32-byte lines each. The Tag Array contains TLB pointers indicating the pages to which the data lines belong and is organized identically to the Data Array. The Data Select Array contains subsets of the primary tag virtual addresses and is used to speed the selection decision between primary and alternate halves. These three arrays constitute the Operand HSB. The TLB Array contains virtual-to-absolute address translations and related address space identification.

Selected bits from the operand virtual address are used to index into all four arrays simultaneously. Both the primary and alternate lines from the Data Array are accessed. The Data Resident Match indicates the requested line is present if both the operand virtual address and the tag virtual addresses point to the same TLB entry, and if the operand virtual address and address space identifier equal the TLB virtual address and STO. A Data Select Match indicates the requested data is in the primary half of the Data Array if the access operand virtual address equals the virtual address in the Data Select Array, but this indication is valid only when the Data Resident Match indicates that the requested line is present.

After the four array accesses and two comparisons, five data items are sent to the OWR complex: (1) the primary data line, (2) the alternate data line, (3) the line present/missing indication, (4) the primary/alternate select indication, and (5) the LJR information. The OWR complex must then select the correct line if present, and rotate, justify, and truncate it according to the LJR information.

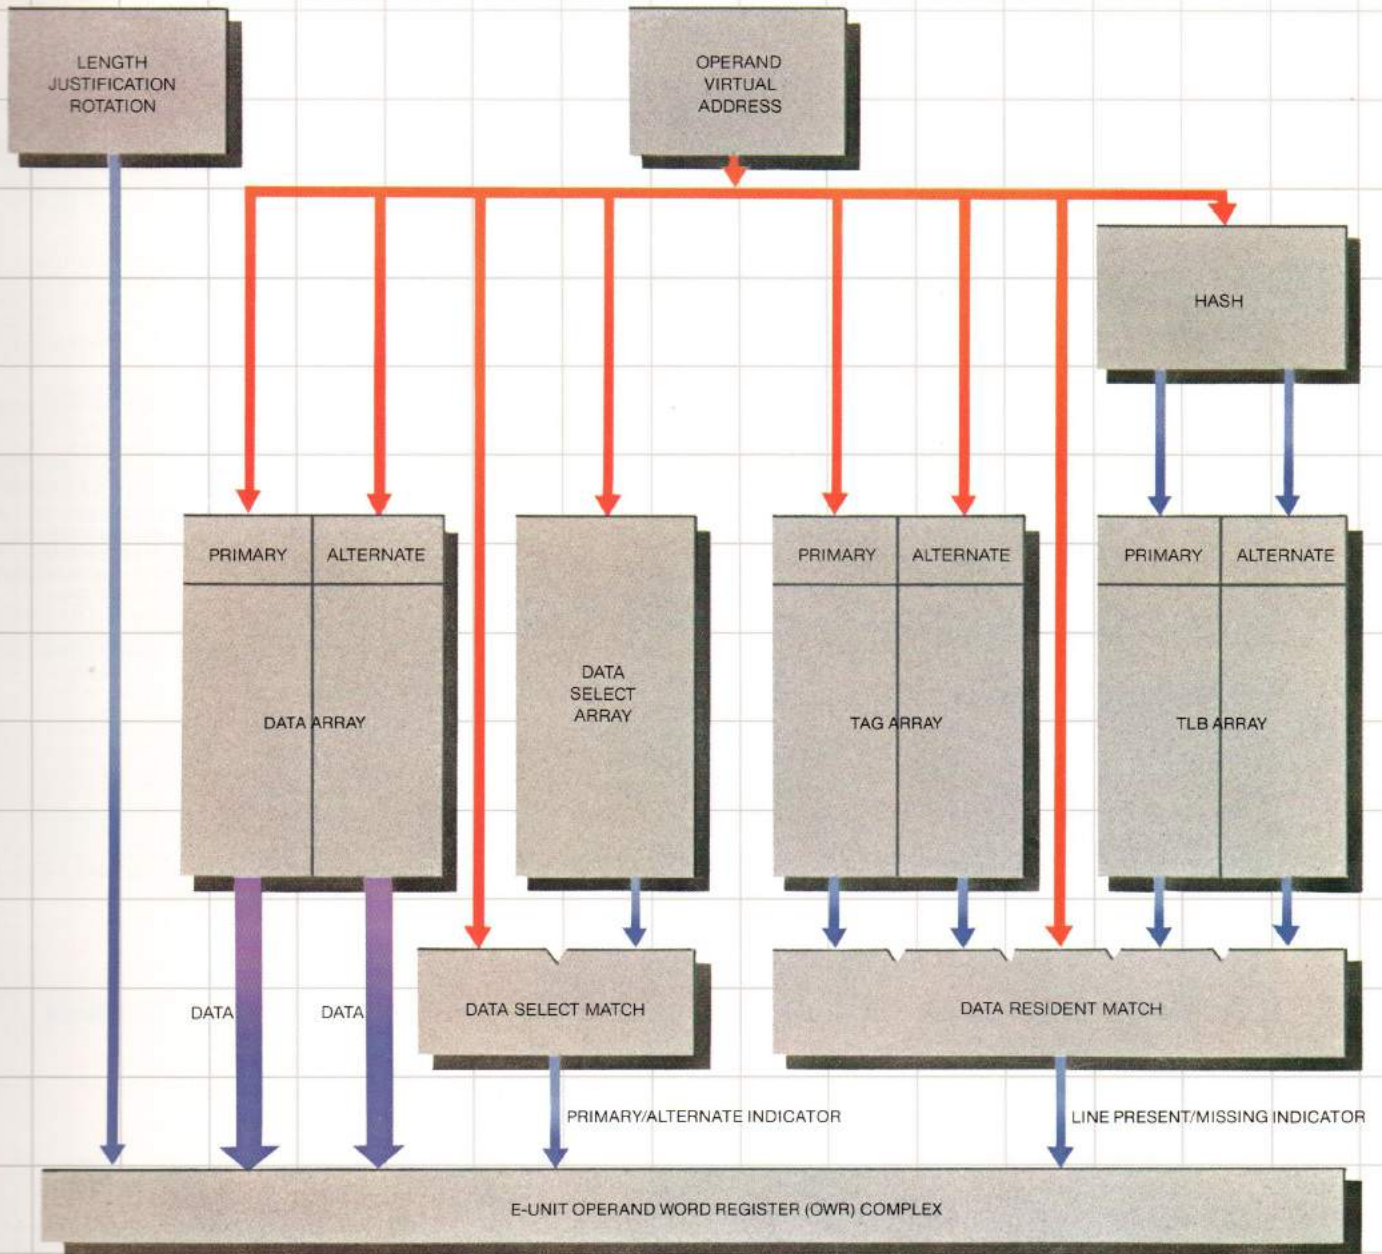
The I-Buffer access is identical to the O-Buffer access except: (1) a special register containing the translation for the current executing page is used instead of the TLB to avoid TLB access conflicts with the O-Buffer access process; (2) different data, tag, and data select arrays are used; and (3) the five data items are sent to the I-Fetch mechanism instead of the OWR.



S-Unit

- Dual 32K HSBs permit separate instruction and operand fetch pipelines.
- Main Memory ECC permits correction of single-bit errors and detection

of all double-bit errors during move-ins.



STORAGE UNIT (S-UNIT) AND OPERAND HIGH-SPEED BUFFER (HSB)

Execution Unit (E-Unit)

The E-Unit provides computational facilities to implement the Amdahl 580's instruction set. It receives data from and returns data to the Operand Buffer (O-Buffer) or the I-Unit Register Facility as appropriate for the opcode.

Performance

In order to keep the pipeline full, the E-Unit Logical Unit and Checker (LUCK) and Execution cycle logic can work on two separate instructions concurrently. Pipeline delays are further reduced by a bypass structure which makes the results of an instruction available for use as an operand in the succeeding instruction through much shorter and faster data paths from the Result Register to the Operand Word Register (OWR) complex.

The E-Unit also incorporates features which reduce the average execution time of an instruction. Primary data paths are eight bytes wide, increasing the amount of data that can be manipulated in a cycle. Particular attention has been paid to optimizing algorithms and logic for instructions which, in typical customer environments, use the most cycles due to their frequency of occurrence and their complexity. For example, the Amdahl 580's E-Unit can perform a decimal addition or subtraction in as little as one cycle by utilizing its three-port, double word Adder and Decimal Correction logic.

Compatibility

The Amdahl 580's E-Unit is micro-coded for flexibility to maintain future compatibility. Microcode is located on the same Multiple Chip Carrier (MCC) as the E-Unit logic, with a structure tailored to E-Unit control. The unique microcode branch technique avoids microcode access delays.

Availability

The E-Unit utilizes residue and parity checks to ensure that computations are correct.

The E-Unit data flow is divided into two parts, corresponding to the two E-Unit phases: the LUCK phase and the Execution phase.

During the LUCK phase, operands arriving from the I-Unit and S-Unit are selected, rotated, justified and truncated, and latched into the OWR. The data then flows through one of three subunits: the LUCK itself, where ANDs, ORs, EXCLUSIVE ORs, logical and sign comparisons, and certain data manipulation functions are performed; the Byte Adder, where floating point exponent addition and subtraction are performed; and the Byte Mover, where functions associated with the Edit, Move Zones, Move Numerics, and Unpack instructions are performed. The results of this phase are then latched into the Staging Platform Registers, ready for use in the Execution phase. During the LUCK phase, early condition code setting is performed to allow more efficient branching.

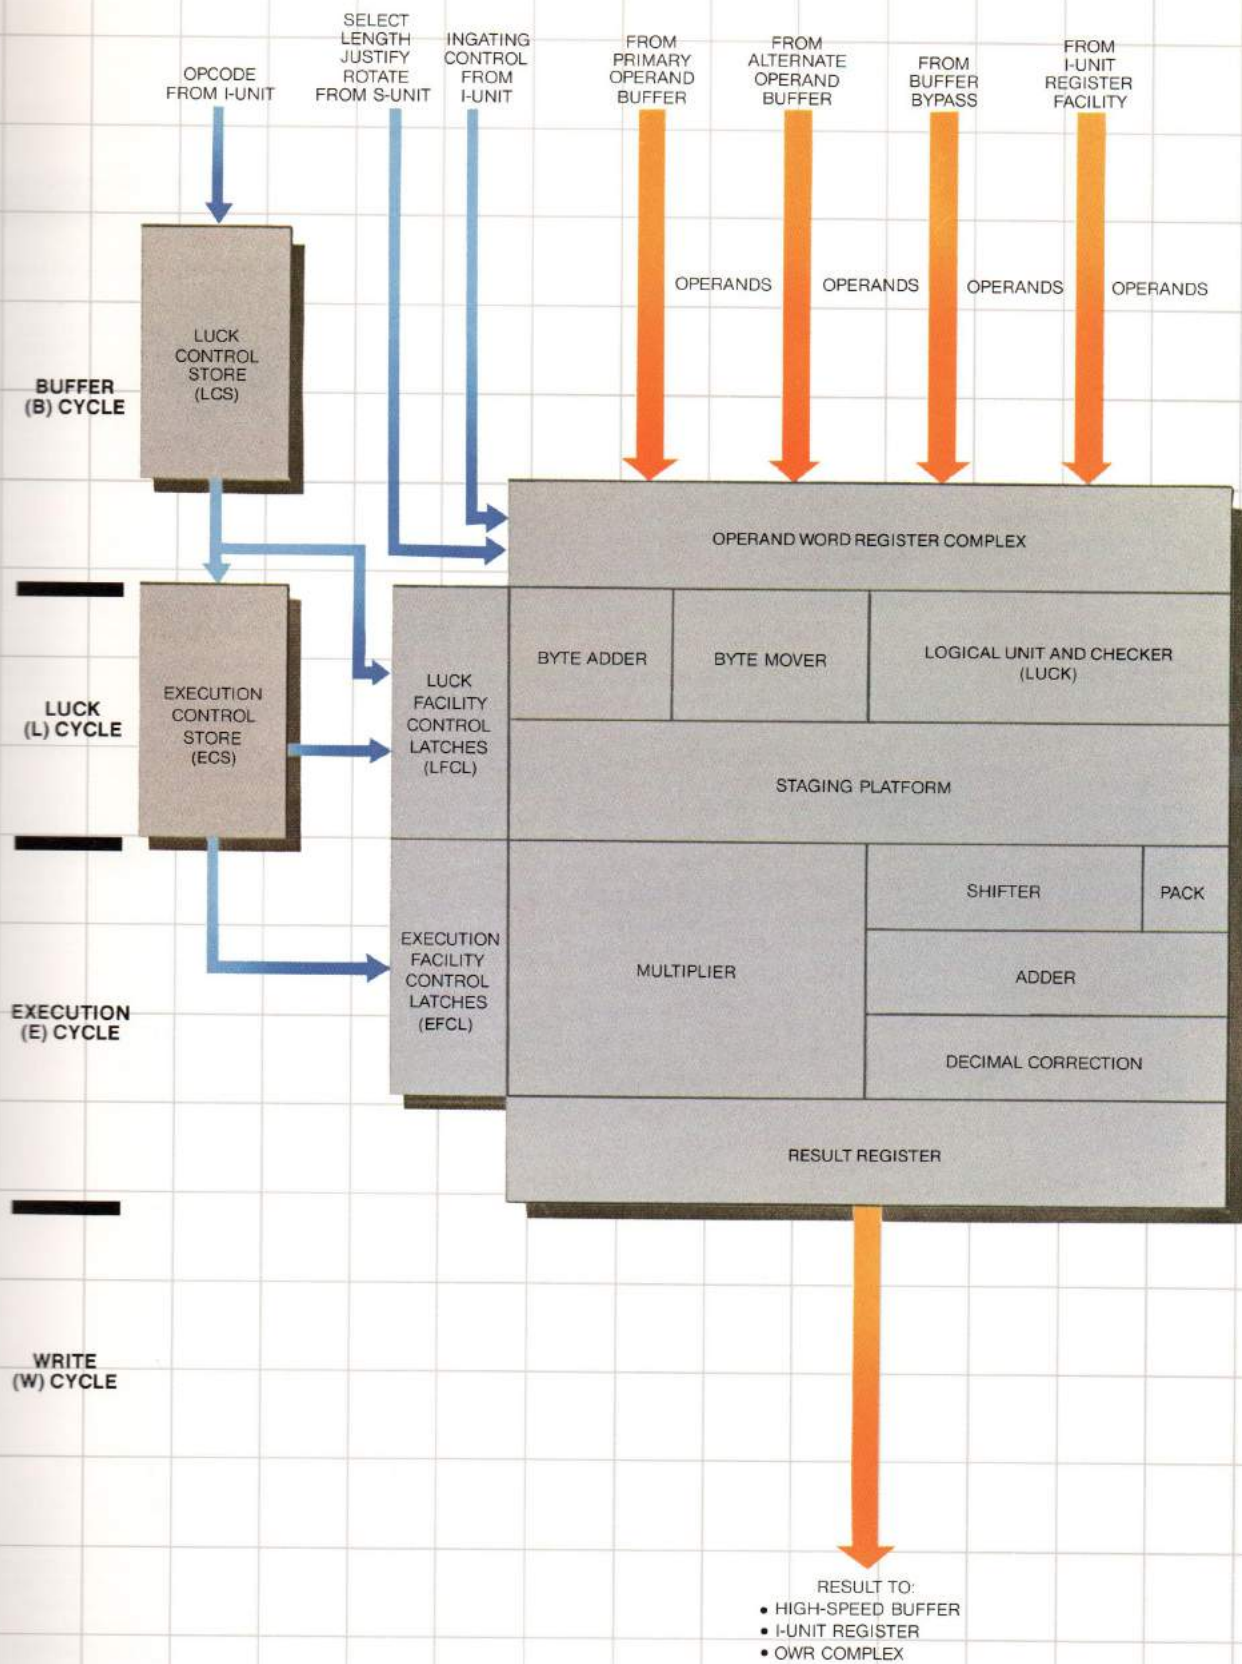
During the Execution phase, data from the Staging Platform flows through one or more of the following units: High-Speed Multiplier; Shifter; three-port, carry propagate Adder; Pack Function; and Decimal Correction Function. The results are then latched into the Result Register.

There are three separate sets of controls for the E-Unit data flow: (1) the I-Unit pipeline controls and LJR information, together with the S-Unit-generated Data Resident Match and Data Select Match indicators, control the OWR ingating; (2) during the Buffer Cycle (B-Cycle), a microcode control word is accessed from the LUCK Control Store (LCS) using the I-Unit-supplied opcode as an index and is gated into the LUCK Facility Control Latches (LFCL), where it controls the LUCK phase data flow; (3) during the LUCK Cycle (L-Cycle), a microcode control word is accessed from the Execution Control Store (ECS) using a pointer from the LCS control word as an index and is gated into the Execution Facility Control Latches (EFCL), where it controls the Execution cycle data flow. Operations requiring more than one cycle in the Execution phase are controlled by chained microcode control words.



E-Unit

- LUCK and E-Unit cycle logic can work on two separate instructions concurrently.
- Primary data paths are widened to eight bytes to allow more processing per machine cycle.
- Algorithm improvements result in significant increases in average instruction rate.
- Distributed Microcode provides flexibility.



Memory Bus Controller (MBC) and Main Storage Unit (MSU)

The MBC provides communication paths and message traffic control between major Amdahl 580 components using the A-Bus and B-Bus. A Data Integrity unit assures that the current version of a data line is accessed where multiple copies of a line can exist in the Operand Buffer (O-Buffer), Instruction Buffer (I-Buffer), Input/Output Processors (IOPs), Console, or Main Memory. An Interrupt Router unit receives interrupts from units external to the CPU and routes them to the CPU. A Timer Complex unit is provided to implement the System/370 timing functions, including the Time-Of-Day (TOD) Clock, TOD Clock Comparator, CPU Timer, and Interval Timer. An I/O Router unit translates logical channel addresses into physical channel addresses and properly formats them for the IOP or Console as appropriate. This translation process allows the user to perform limited reconfiguration of physical channels without requiring a software change. A Bypass unit provides a path for routing A-Bus messages onto the B-Bus.

The Main Storage Controller (MSC) receives data requests from other Amdahl 580 units and generates control and timing signals required by the Main Storage Unit (MSU) to honor these requests. The MSC generates Error Checking and Correction (ECC) for all data stored in the MSU. Error detection and correction are performed when data and these ECC codes are read from the MSU.

The MSU contains data and key storage arrays for up to 32 megabytes of Main Memory and honors requests for 32-byte line move-ins and move-outs.

Performance

With four-way line interleaving and four-way quarterline multiplexing, the MSU provides up to 32 megabytes of memory with a fast basic access for 32-byte lines. Double word data bus paths transfer eight-byte messages plus byte parity information between the MBC and functional units every cycle. The Amdahl 580 bus system is optimized for Main Memory data fetches which are the most common bus transactions.

Compatibility

In many cases, I/O channels can be reconfigured without System Generation by utilizing the I/O Router function.

Availability

Histories of bus transactions and memory errors are maintained to aid in fault detection. Error Checking and Correction (ECC) is performed on a quarterline basis for reliability.

The A-Bus is logically treated as a single data path, but is physically composed of three data paths: the CPU A-Bus, IOP A-Bus, and Console A-Bus.

The most common MBC message is a Main Memory read request from the S-Unit, an IOP or Console, which arrives at the MBC on the A-Bus. The MBC passes the message to the Main Storage Controller (MSC) unit, which uses the message opcode and address portion to create control signals for the Main Storage Unit (MSU). The MSU array accesses the four quarterlines from one of the four interleaves and latches them in the Main Storage Data-Out Register (MSDOR). The quarterlines are sent to the S-Unit, an IOP, or Console over the Move-in Data Path to complete the access.

To store a line of data, the CPU or an IOP sends a store message consisting of a header segment containing the line address followed by four data segments each containing a quarterline. Console operation is similar but restricted to one quarterline. When the header segment arrives at the MBC on the A-Bus, it is passed to the MSC for control purposes. A data segment containing a quarterline arrives at the MBC on each of the next four cycles; ECC codes are generated and are latched into the Main Storage Data-In Register (MSDIR). They are then stored into one of the four main storage array interleaves.

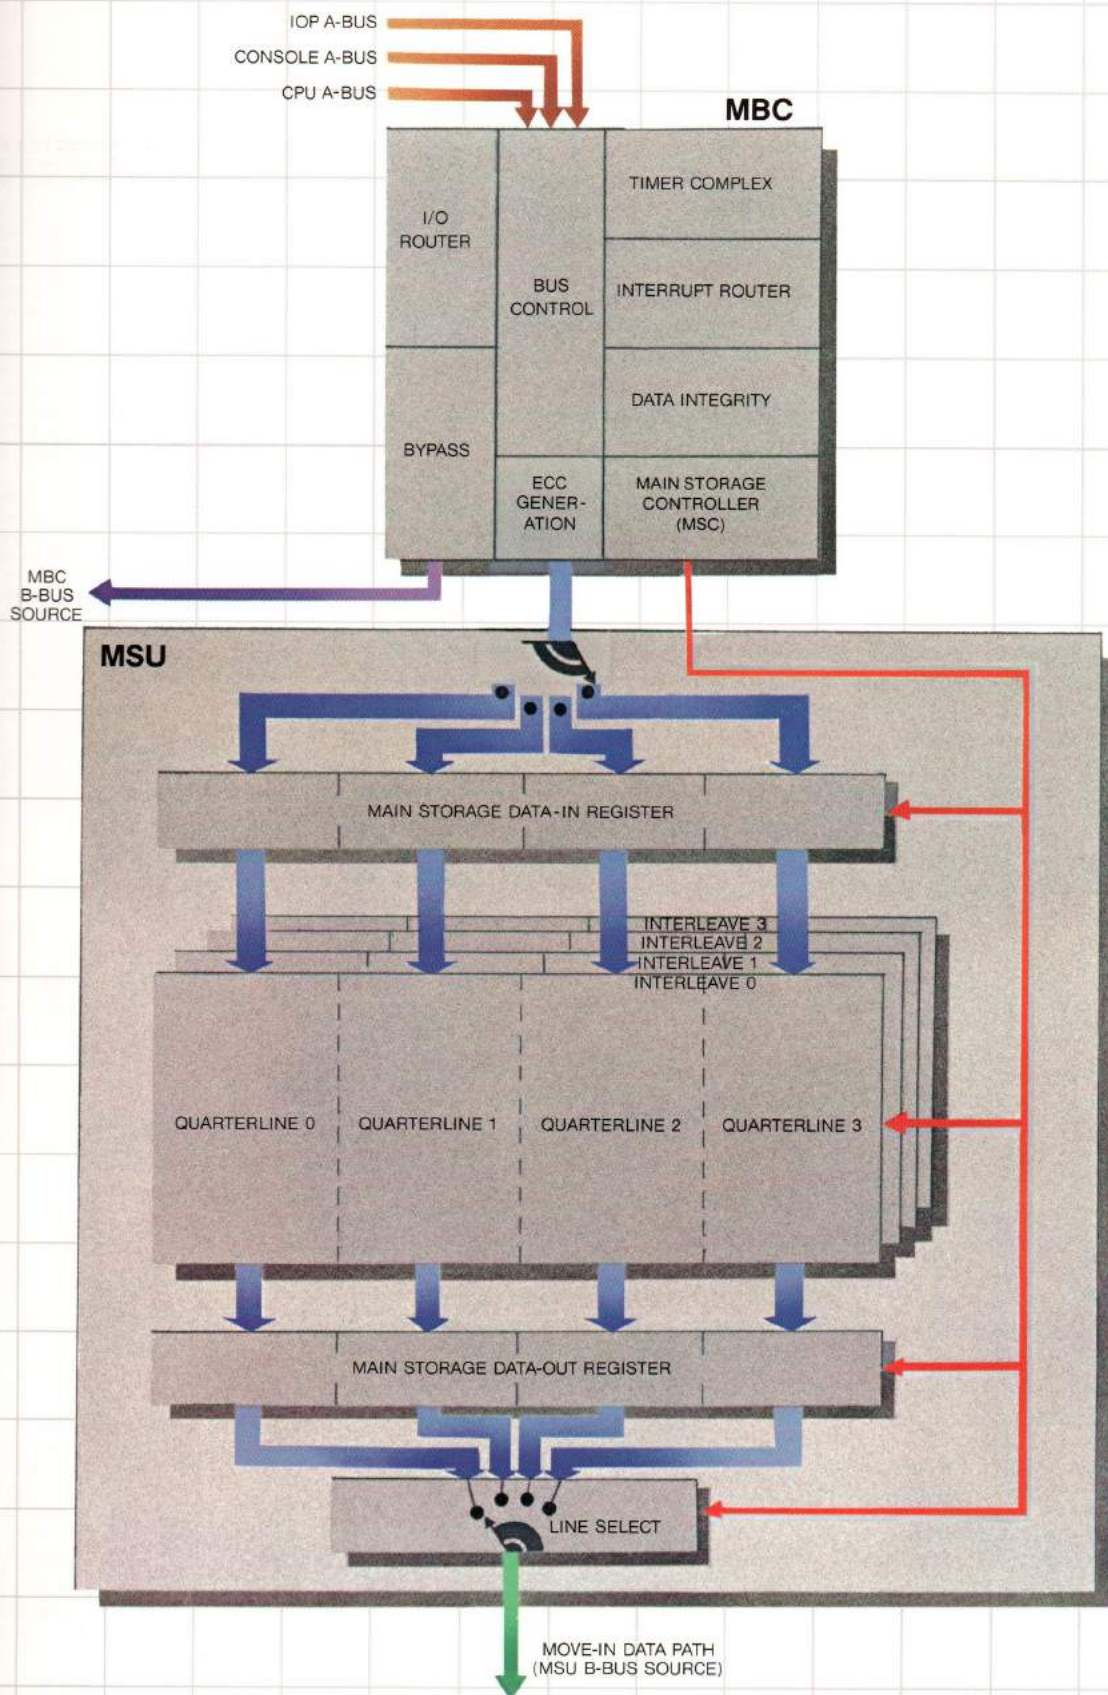
I/O requests from the CPU arrive at the MBC on the A-Bus and are passed to the I/O Router unit which, in turn, sends them over the B-Bus to the IOP or Console Processor as appropriate. I/O interrupt messages from the IOPs and Console arrive at the MBC on the A-Bus and are passed to the Interrupt Router unit which then sends them to the CPU over the B-Bus.



MBC/MSU

- Four-way interleaving and four-way quarterline multiplexing provide fast access to up to 32 megabytes of Main Memory
- The I/O Router

simplifies channel configuration.



MEMORY BUS CONTROLLER (MBC) AND MAIN STORAGE UNIT (MSU)

Input/Output Processor (IOP)

The basic Amdahl 580 includes one IOP which provides 16 Block Multiplexer Channels, and an optional second IOP adds up to 16 Block Multiplexer Channels. Each Block Multiplexer Channel has 256 subchannels.

An Amdahl 580 IOP is the primary interface between peripheral devices and the CPU. An IOP receives I/O commands from the CPU via the bus system and uses the bus system to handle data and command fetching, data and status storing, and channel interrupts. An IOP's Interface Handlers perform channel bus and tag manipulations and data buffering to provide the external interface to the user's peripheral devices.

An IOP consists of three functional units: The I/O Controller (IOC), the Bus Handler, and the 16 Interface Handlers. The IOC and Bus Handler are shared among the 16 Block Multiplexer Channels, and are implemented on the IOP MCC. Separate Interface Handlers are provided for each channel and are implemented on special cards in the channel frame portion of a main-frame using the same LSI chips and RAMs as are used in the stack.

The IOC provides an IOP's processing power and controls its Bus Handler and 16 Interface Handlers.

The Bus Handler provides the interface and buffering between an IOP and the other Amdahl 580 functional units.

The Interface Handlers perform normal data transfer operations, including channel bus and tag manipulation and data buffering.

Performance

Maximum data rates for individual channels and the aggregate of all channels are significantly improved by the Amdahl 580's IOP design. Individual channel bandwidth can be as great as six megabytes per second, and aggregate bandwidth is approximately 50 megabytes per second.

Actual data rates depend upon the control unit interface and channel protocol. This performance is achieved by reducing both inter-channel interference and interference between the I/O subsystem and the CPU.

Interchannel interference is reduced by the IOP's major shared processing component, the IOC, and by the use of separate Channel Interface Handlers. Because the IOC is a very powerful time-sliced, microcoded processor, all the channels can concurrently execute the same microprogram or completely independent microprograms. Each channel's Interface Handler is solely responsible for its own data buffering and bus and tag manipulation.

Interference between the I/O subsystem and the CPU is reduced by using the Amdahl 580's bus system to fetch commands and data directly from Main Memory instead of through a shared High-Speed Buffer.

A new Amdahl 580 feature, Subchannel Queuing, holds I/O requests that have been rejected by a busy channel or device, permitting efficient restart after that channel or device becomes available. The CPU load associated with restarting the rejected I/O request is significantly reduced.

Compatibility

The Amdahl 580's IOPs are compatible with System/370 I/O architecture for Block Multiplexer Channels. To provide flexibility for response to future extensions, the IOC is implemented as a microcoded processor, and the Interface Handlers can be adapted to new channel protocols with little or no effect on the IOC or Bus Handler.

Availability

The IOP's modular design and interface simplicity reduce complication and thereby improve both reliability and serviceability. The implementation of all IOP components in LSI further enhances reliability.

The Input/Output Controller (IOC), Bus Handler and 16 Interface Handlers work together to accomplish I/O functions.

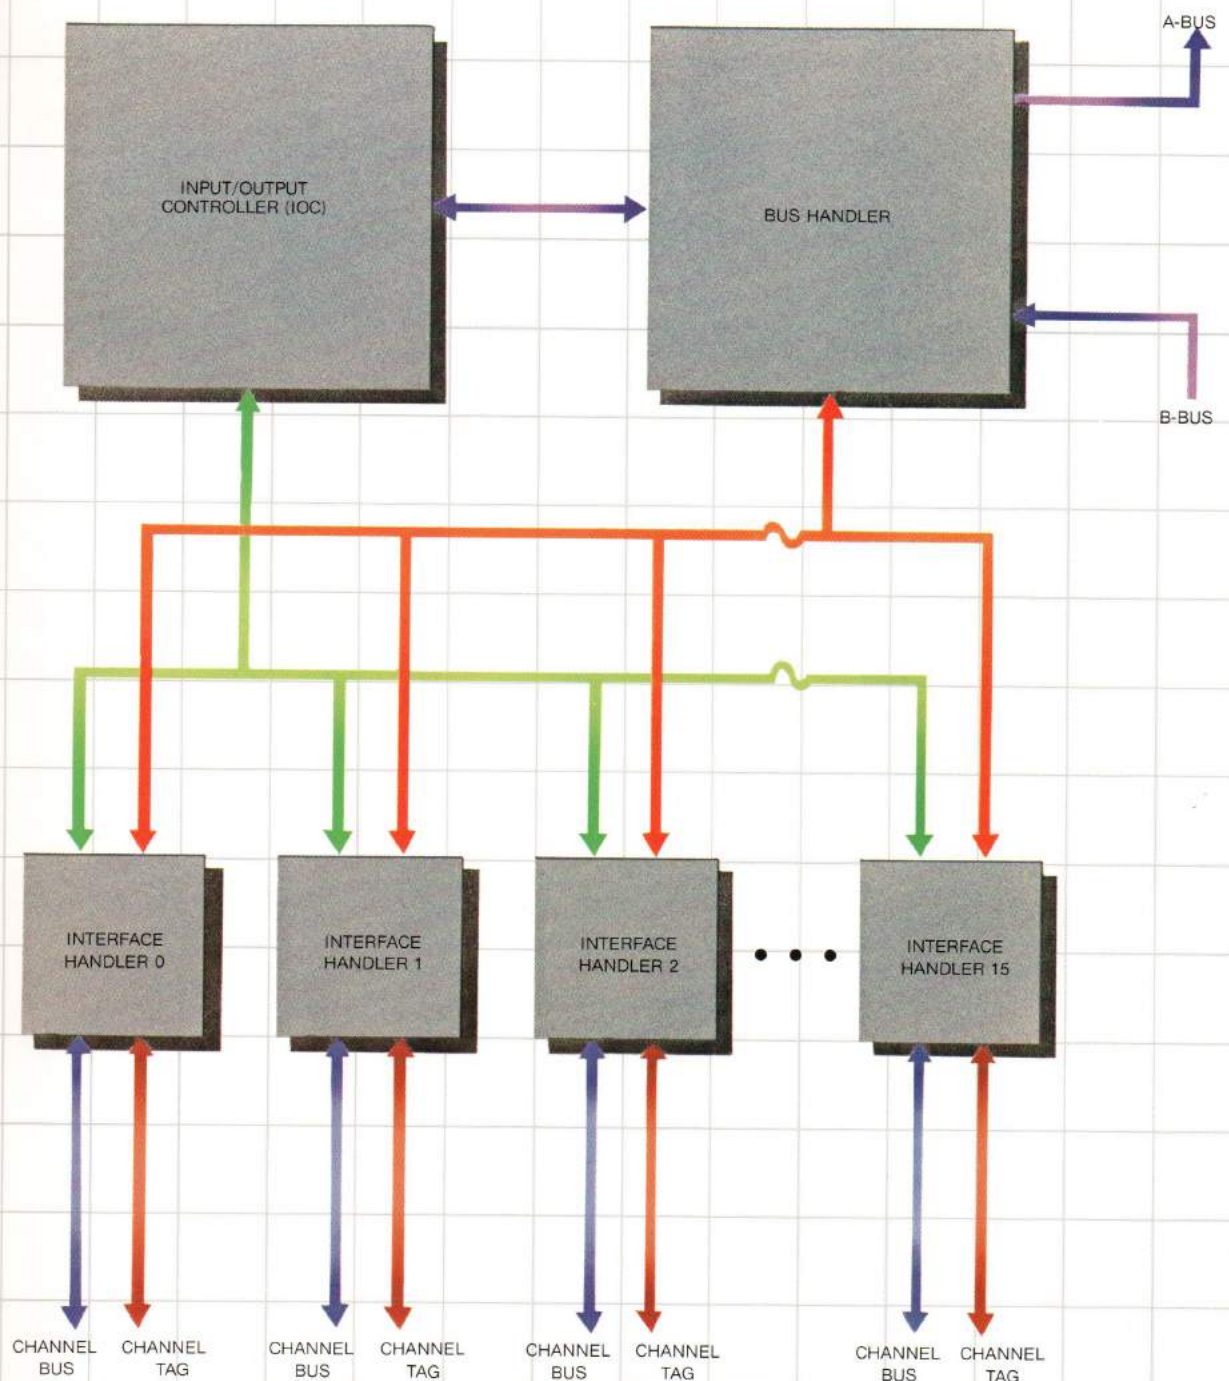
A Main-Memory-to-device or device-to-Main-Memory data transfer specified in a Channel Command Word (CCW) is initiated by the IOC. It is then controlled by the Bus Handler and associated Interface Handler. During a device read, the Interface Handler accepts data from the device, buffers it, and then sends it to the Bus Handler. During a device write, the Interface Handler accepts data from the Bus Handler, buffers it, and then sends it to the device. IOC intervention is minimal, required only to continue data transfers with lengths greater than the Interface Handler's buffer capacity or for termination.

The Bus Handler assembles data line fetch and store request messages for the Interface Handlers. It also assembles command line fetch and store and interrupt messages for the IOC. These messages are transferred to other Amdahl 580 functional units via the A-Bus. The Bus Handler also receives unsolicited request messages such as Start I/O requests. Solicited reply messages containing IOP requested CCW lines and data lines from the other Amdahl 580 functional units are received via the B-Bus and disassembled for transfer to the Interface Handlers and IOC. B-Bus monitoring by the Bus Handler occurs independently of the IOC. During data transfer, data fetch and data store messages are sent independently by the Bus Handler.



IOP

- Individual channel data rates of up to six megabytes per second and an aggregate data rate of approximately 50 megabytes per second balance CPU performance.
- Subchannel Queuing permits efficient restart of rejected I/O requests and reduces CPU load.
- All IOP components are implemented in LSI.



INPUT/OUTPUT PROCESSOR (IOP)

Console Complex

The Amdahl 580 Console Complex provides a System/370-compatible operator's console interface and the additional control, diagnostic, and measurement functions that are required by the Amdahl 580's architectural extensions. These extensions include AMDAC, which enables a user's Amdahl 580 to be diagnosed via telephone line from anywhere in the world.

The Console Complex also includes the Amdahl 580's Byte Multiplexer Channels. The Interface Handlers for the two Byte Multiplexer Channels are very similar to and are located with the Input/Output Processor (IOP) Interface Handlers.

The Console Complex consists of: (1) a microcoded CPU capable of executing a significant subset of Amdahl 580 instructions and functions; (2) two megabytes of memory; (3) an I/O channel; and peripheral devices which include: (4) a hard disk; (5) two floppy disks; (6) up to two local and two remote 3277-like CRT/keyboard stations; (7) a modem controller for AMDAC; (8) an Amdahl 580 diagnostic latch scanner; and (9) a microcomputer-based support processor which functions as the Console's own console.

Performance

The Console CPU, channel, control units, and the two Amdahl 580 Byte Multiplexer Channels are implemented as time slices on a separate and special version of the I/O Controller (IOC) used in the IOP. This microprogrammed, time-sliced processor, implemented on its own MCC, is teamed with special microcode and interfaces to give the

Console the computational capability of many medium-sized computer systems, and to give the Byte Multiplexer Channels bandwidths of .2 MB/second. By grouping the slower Byte Multiplexer Channels with the Console, the IOPs are free to handle faster Block Multiplexer operations exclusively. The Console's peripheral devices act as integrated control units using shared memory instead of bus and tag cables for communication and thus provide more efficient data transfer.

Compatibility

Operator interface to the Amdahl 580 Console is System/370-compatible. Maintenance functions are compatible with Amdahl 470 system maintenance procedures. For adaptability to future change, the Amdahl 580 Console Processor is implemented in microcode, and the Console control program is written in an Amdahl 580 instruction subset.

Availability

The Console Complex is Amdahl's primary tool for performing both local and remote system diagnosis. In the event of a failure, its extensive capabilities facilitate the rapid return of the Amdahl 580 to service.

The Console scan facility enables the Amdahl field engineer to examine the status of almost any Amdahl 580 latch and to change most of them. Any function which can be performed at the Amdahl 580 installation can also be performed or monitored via telephone line by an Amdahl specialist at one of the Amdahl Diagnostic Assistance Center (AMDAC) locations around the world.

The Console Processor's reliability is enhanced by its all-LSI implementation on a single MCC. Serviceability is enhanced because the Console Processor has its own microcomputer-based support processor to assist in diagnosis on those rare occasions when the Console Processor itself fails.

The Console Processor, the Console peripheral device control units, and most of the logic for the two system Byte Multiplexer Channels are implemented on a special version of the very powerful, time-sliced microcode processor found in the IOP.

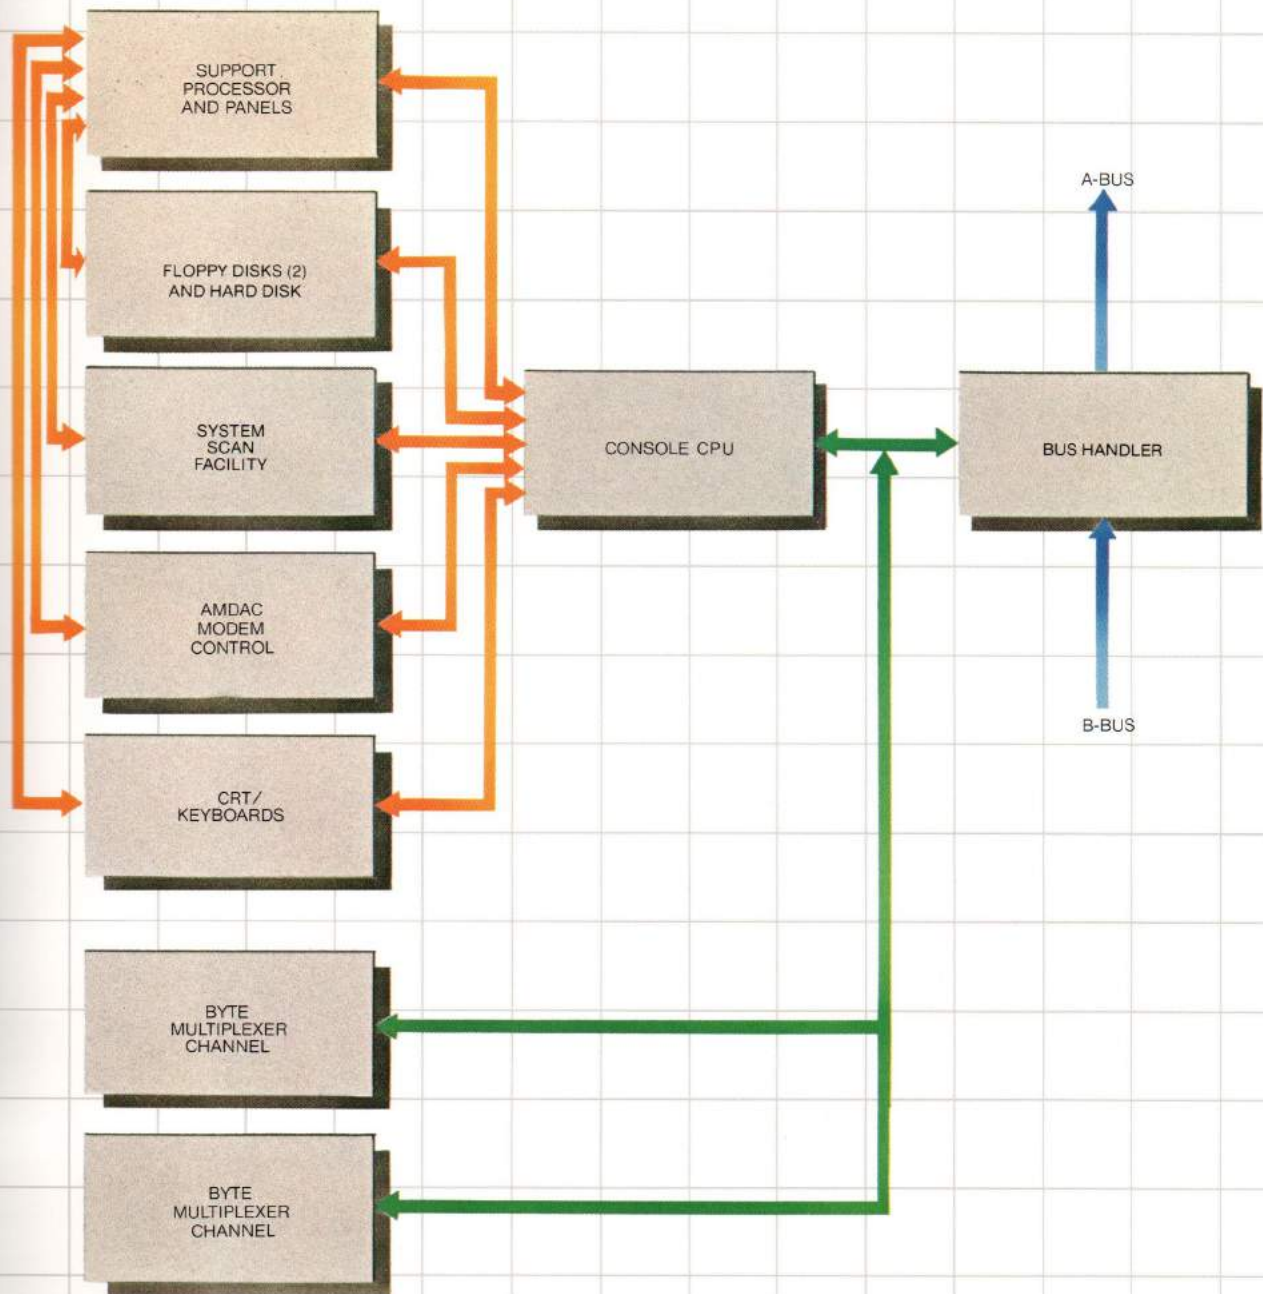
The Console Bus Handler ensures that messages to and from other Amdahl 580 functional units are routed correctly for the Console units.

The Console Processor and Bus Handler are housed on an MCC in the stack. The Interface Handlers for the two Byte Multiplexer Channels, certain support functions, the Support Processor, and other console peripherals are housed in the channel frame.



Console

- LSI implementation on a single MCC enhances Console Processor reliability.
- AMDAC remote diagnostic interface and system latch scanner facilitate serviceability.
- Console programming uses a subset of the Amdahl 580 instruction set to enhance flexibility to retain compatibility.



CONSOLE COMPLEX

The many innovations and features of the Amdahl 580 can be categorized according to their impact on the principal benefits of Performance, Compatibility, and Availability.

PERFORMANCE

CPU performance of two times the 470V/8 in typical commercial environments

- Average cycles per instruction reduced significantly
- Pipeline improvements
 - Instruction-per-cycle maximum execution rate
 - Short five-phase pipeline
 - Early condition code setting
 - Reduced interlocks
 - Separate Instruction Fetch and Instruction Execution
 - LUCK and E-Cycle logic can work on different instructions concurrently
- Primary data paths widened to handle double words
- Dual 32K HSBs: one for instructions and one for operands
- Algorithm improvements in I-Unit and E-Unit
- Four-way line interleaving and four-way quarterline multiplexing for fast access to up to 32 megabytes of Main Memory
- Cycle time improvement
 - Technology
 - 400 picosecond gate delay
 - 400 circuits per chip
 - New fast RAMs
 - Packaging
 - 121 chips per MCC
 - RAMs and logic on same MCCs
 - Eight MCCs for basic system
 - Stack houses MCCs in 5.6 cubic feet and incorporates buses in its walls

Channel performance improvements

- Up to 32 Block Multiplexer and two Byte Multiplexer Channels
- Aggregate data rate of 50 megabytes per second
- Individual channel data rates of up to six megabytes per second
- 256 subchannels on every channel
- Reduced interference
- Subchannel Queuing

COMPATIBILITY

IBM large-scale systems and Amdahl 470 compatible

Flexibility to adapt to future change

- Distributed Microcode in I-Unit, E-Unit, IOP and Console
- Macrocode, a new class of firmware
- Channel Interface Handler design adaptable to new protocols
- Console programming uses subset of Amdahl 580 instruction set

AVAILABILITY

Reliability

- Technology and packaging
 - Fewer components
 - MCC interconnection through printed circuit boards
 - IOPs and most of Console implemented in LSI
- Error detection and correction
 - Main Memory ECC
 - Bus parity checking
 - E-Unit parity and residue checking
 - Instruction retry

Serviceability

- Fault isolation
 - Function per MCC organization
 - Fault isolation circuitry
- History RAMs for microcode and bus transaction tracing
- AMDAC and system latch scanner
- Microcomputer console support processor